Structuring for Strategic Cyber Defense: A Cyber Manhattan Project Blueprint

O. Sami Saydjari
Cyber Defense Agency, LLC
ssaydjari@CyberDefenseAgency.com

Abstract

In February 2002, more than 50 leaders in the information assurance field warned the President of the United States of a national strategic vulnerability in the country's information infrastructure that could cause mortal damage. Six years later, some motion in the direction of a government strategic investment is beginning to get under way. This essay will address the key capabilities needed at a national scale and how those capabilities might drive a vigorous research and technology agenda. The text also addresses several imperative questions: How might we organize a government activity in which many agencies surely need to be involved yet must march in a coherent direction? What lessons can we learn from the post-Sputnik era to regain leadership in the space race? Has a cyber Sputnik already launched, and, if so, is the US already behind in the cyber space race?

1. Introduction

Like many industrialized countries, the US is vulnerable to a strategically crippling cyber attack from nation-state-class adversaries. Most of our real-world critical assets are controlled via cyber space and large primary value such as intellectual property is contained exclusively in cyber space, which makes it as legitimate a part of our territory as physical land, thus the government must "*provide for the common defense"* of this new territory as prescribed in the preamble of the US Constitution [1]. A strategic multi-billion-dollar investment run by the country's best experts can mitigate our risk of a cyber attack if we start now [2].

This essay addresses the question of why such a program is needed, what shape it might take, what sorts of things would need to be done, how we might organize to accomplish the Herculean tasks, and how we might get started today on such a path.

2. National Strategic Threat

The Threat: What is the problem, and what is the solution? For the problem, we must ask if a strategic national vulnerability exists, what its scope is, and how bad "bad" can get. Without understanding the detailed nature of the problem, the efficacy of any proposed strategy is unknown. We must also ask how a proposed national strategy would solve the problem, and what happens if it doesn't. These seem like childishly simple questions, but the answers have been elusive. Indications are that national economic devastation is quite possible, and when we're in the middle of a disaster isn't the time to start thinking about how to respond. Preparing for cyber war will take several years and require infrastructure instrumentation for critical computer systems, experienced cadres of defenders who are well trained and exercised (through simulation, for example), control systems to execute strategic responses (by orchestrating a large scale of configuration changes in myriad security devices such as firewalls, for example), effective architectures to mitigate risk (such as by layering as discussed below), and a national program to engineer defensive capabilities into our critical infrastructure and to create new capabilities through a vigorous research investment. Thus, understanding the problem is an immediate need.

Planning. To help the government better understand the gravity and nature of the threat, a nonprofit cyber policy advocacy group known as the Professionals for Cyber Defense (PCD, see www.uspcd.org) created a strategic attack scenario dubbed Dark Angel. Our small planning team included a campaign planner, two experts focused on the financial sector, three in electrical power, and one in transportation. We assumed only unclassified critical

1063-9527/08 $25.00 © 2008 IEEE
DOI 10.1109/ACSAC.2008.53

IEEE computer society

infrastructure vulnerabilities. Our intent was to illustrate the strategic damage a robust campaign that used multiple attack paths could cause and to create a plan with sufficient detail to convince experts in the domain of the plan's validity. The plan took roughly 30 days to create. We assumed the adversary had three years of preparation, $500 million, and 30 days to actually execute the attack. The attack campaign's goal was to destabilize the US and depress the economy with attacks on critical infrastructure, thus reducing our ability to project military power, depleting our will to fight, and creating panic and distrust in the government.

Our strategic campaign objectives included crippling rail transportation, rupturing oil and gas pipelines with improper control (for example, with cyber attacks similar to the one on the Soviet Trans-Siberian pipeline that caused a three kiloton explosion, as described in "At the Abyss" by Thomas Reed), and creating widespread power outages by destroying hard-to-replace generators and power-line transformers with improper computer control commands. We also simulated attacks on financial services sectors to create mass confusion in transaction settlement systems, flooded 911 emergency-response systems with computer-controlled false alarms to create widespread panic, and disabled Internet service by performing denial-of-service attacks on the 13 primary Domain Name Servers (as has been partially done in actual cyber attacks).

In our simulated campaign, we spoofed attack attribution to focus attention in the wrong direction, launched lethal first strikes (for example, hitting first responders and backups before hitting primary cyber targets), used a rolling attack barrage to interfere with recovery processes, delayed instrument attacks until they were no longer needed in the campaign, bought cyber mercenaries and insiders as needed to gain capabilities and access, launched non-cyber (physical) attacks on "tough" targets as needed, employed psychological operations to create distrust in infrastructure and manipulate public opinion, and hampered the military by disrupting civilian re-supply chains.

We vetted our simulated attacks with experts in each of the key critical infrastructure domains and verified the essence of the plan and its likely effects. There was some uncertainty about the consequences of some attacks—even now—but this was due to a lack of knowledge among the entire community to fully assess such consequences. It would be hubris to think our adversaries don't already have a plan in place that's substantially better than our brief sketch or that their capabilities to execute such an attack aren't improving.

Follow-on. A proper national strategic threat assessment would parallel that of Dark Angel and involve top industry experts and business leaders, and military campaign planners, economists, policy makers, and others as needed. Sharing across industry should be encouraged and rewarded. From a management perspective, the assessment should carry presidential authority and priority. Essentially, the task should involve three separate teams: one for planning and completing a concrete plan, one to execute the plan to the extent needed for demonstration purposes, and one to review the results for validity.

The assessment must start from the premise built into Dark Angel: that cyber warfare is equivalent to economic and social warfare. Diagnosis of vulnerability sources must reflect that the organization and design of our production systems is often more important than cyber defense technology in determining the destruction's nature and extent. What to defend and what kinds of damages to prevent aren't self-evident without such an assessment.

For illustrative purposes, we estimated the resources needed to analyze six critical infrastructure domains would take roughly $70 million, 300 top-talent experts, and 9 calendar months. The final report would be a definitive estimate of our true national strategic vulnerability to cyber attacks, a compelling case for action, and the basis of a prioritized program plan.

Although we originally created the scenario to help policy-makers understand the gravity of the strategic threat, further thought suggests that such scenarios can provide invaluable driving metrics for assessing mitigation programs of the type discussed in this paper. An excellent and eerily parallel analysis for the biological threat in terms of such a scenario appears in a paper by Richard Danzig [3].

3. A New Manhattan Project

As part of a dialogue with the government in 2002, the PCD elaborated on the proper solution to the strategic vulnerability sketched out by our Dark Angel analysis. Cyber war defense requires orders of magnitude more government involvement and resources to avoid overwhelming national damage from strategic attacks. The PCD recommended that the government (1) step up to a strong defense role against serious attacks, (2) focus on countering strategic attacks that have real-world effects, (3) develop a top-down architecture and engineered approach to the defined problem, (4) acknowledge that current technology is insufficient to defend against cyber war, and (5) divide the cost burden between the owner (to protect critical private cyber assets) and the government (to protect the integrity of the national commons).

As mentioned earlier, the PCD chose the name "Cyber Manhattan Project" to reflect the gravity, urgency, priority, focus, top-talent, and funding levels needed. We acknowledge that aspects of the analogy are inapt, such as the fact that (1) there is no single, easily measurable artifact (such as a bomb), (2) a broad spectrum of talent and organizations must be involved, (3) much of the work must be conducted without classification constraint, and (4) once an initial capability is achieved, a continued investment will be needed to maintain our cyber defense's effectiveness. We sketch the program below.

Vision. We must rapidly overcome our nation's vulnerability to coordinated strategic cyber attacks from serious enemies.

Urgency. Major potential adversaries are actively pursuing cyber war capabilities, indicating the increasing probability of future cyber campaigns. Moreover, current cyber defenses and best practices are ineffective, active measures to shut down our adversaries' abilities to attack through physical access will drive them to cyber space, and we face potentially greater vulnerability and lethality from combined cyber and physical attacks. Finally, developing a defense to this threat is a multiyear effort, so we can't wait until we're in the midst of our first major strategic attack campaign to start defending against it.

Priority. A major initiative on the order of the Cyber Manhattan Project is the right path to address our current situation. The offensive threat is growing, so defense must be fielded at a faster rate. A top-down approach with a driving architect can address the problem and achieve the requisite objectives, but bottom-up efforts, even if coordinated, leave gaps because there's no ownership of key parts of the problem. Cyber defense mechanisms must integrate into a coordinated system, and cyber defense operations must comprise a fully integrated defensive force. For success, the creation of national cyber defense capabilities must be a national funding priority. Can you imagine the original Manhattan Project succeeding without such a focus?

Feasibility. Not only is the creation of national cyber defense capabilities critically urgent and important, it's also feasible. (1) Technically, many effective defensive technologies exist but are in research stages and must be transitioned to operational use, some already have limited field testing, and others already exist to address broad classes of novel attacks. Moreover, the required computational resources for intensive activities such as correlation of attack and modeling/simulating attack strategies and tactics are available today. Ongoing research sponsored by the National Security Agency (NSA), National Science Foundation (NSF), Department of Defense (DoD), Department of National Intelligence (DNI), Department of Homeland Security (DHS), and others is beginning to address additional hard science problems. (2) Economically, we can make a national business case for investing in a program intended to avoid the expected financial losses from strategic cyber attacks and ensure the proper public–private sharing of the burden. (3) Operationally, we can manage the complex infrastructure though judicious use of automation with a capable cadre of defenders. Through a combination of reasonable fire-code-like cyber security standards, improved operational guidance, and trained/experienced personnel, we would also be able to contain mission and cost impacts in the short term while we develop new capabilities. (4) Politically, public awareness of the threat is likely to make needed investments and standards acceptable. Industry is increasingly aware that nation-state-level attacks are a concern beyond its current ability to handle, yet such attacks threaten business continuity and the economic foundation of the entire country. With proper financial incentives and partnering for workable solutions, industry is likely to openly embrace government involvement and protection. (5) Finally, from a schedule perspective, a phased rollout of capabilities based on threat prioritization and available technologies is also feasible. Success is certainly not assured, but the alternative is to begin radically reducing our dependency on computing systems, which would seriously degrade our national competitiveness and suppress economic growth.

Project Description. We need an aggressive, goal-directed, high-priority, national program to address these types of high-level, far-reaching threats. To do this, we must engage the brightest scientists, business experts, and engineers and provide them with adequate resources. To guide the program with strategic objectives, we need a top-down architecture that establishes concrete cyber defense capabilities on a specific timeline, including near-term capabilities within three years.

The cyber vulnerabilities in our infrastructures have become deeply embedded and widespread through the economic forces that drive individual companies to reduce costs by adopting the most widely available and interoperable technologies. It won't be easy to develop a cyber infrastructure that can resist strategic attacks: it will require short-term actions as well as a long-term plan and a willingness to keep that plan in focus over a number of years.

4. Strategic Cyber Capabilities Needed

In the context of this work, a *capability* is the ability to defend cyber space in some specific fashion. Some cyber defense capabilities to include in a strategic national plan are as follows: (1) system resiliency and quick recovery from partially successful attacks; (2) a national cyber Command, Control, Communication, and Computer Intelligence, Surveillance, and Reconnaissance (C4ISR) system to measure and control mechanisms at multiple echelon levels; (3) a national threat assessment system to drive decisions at some "required" level; (4) cyber firebreak mechanisms and architectures to slow down attacks and reduce potential damage; (5) methods to gather intelligence and inject uncertainty through strategic deception; (6) models and simulations of the enemy, thereby honing our defenses before incurring damaging strategic cyber attacks; and (7) approaches for identifying and understanding available and acceptable responses from technical, strategic, legal, economic, and political perspectives.

As a step toward a security research plan that includes such capabilities, we should identify end-states—goals in terms of how we want our systems to ideally operate. This fresh perspective includes the overall strategic picture and connects clearly with strategic actions that significantly mitigate strategic vulnerabilities. If, for example, the nation has a capability to quickly recover its critical information infrastructure, then the end-state is that strategic attack damages are mitigated and critical services are restored quickly, possibly deterring adversaries from attempting a future attack.

Desired End-States. The National Cyber Defense Initiative (NCDI) Opening Moves Workshop [4] identified important end-states, the outcome of a 10-year research effort to create critical capabilities. The following end-states appear in the workshop proceedings:

- *Continuity of Critical Information Infrastructure Operations.* Create technology that would be the basis for a resilient US cyber infrastructure that would sustain critical functions in the face of attacks, including those that could be affected by determined adversaries.

- *Well-Defended Critical Assets.* Make it economically prohibitive for an adversary to cause strategic damage to critical US infrastructures. Currently, adversaries can attack critical systems without investing substantial resources.

- *Local/Global Cyber Situation Awareness*. Know what's on critical system platforms, what's connected to the network, who's on the network, the traffic flowing over it, and the threats to it. Create cyber early warning systems while maintaining privacy protections for citizens. Move from today's intrusion detection systems that can detect simple previously seen attacks locally to much more effective ones that can see highly sophisticated, novel, covert strategic attacks against information infrastructure.

- *Confidentiality-Preserving Systems.* Prevent unauthorized access and exfiltration of critical information and intellectual property. Ensure accountability for information flows within systems so that information is shared with those intended to have it. Much of the highly valuable information lost from today's systems is "protected" by perimeter devices such as firewalls. New mechanisms and architectures are needed.

- *Extensible Systems that Safely Embrace New Technology*. Confidently add new functions without compromising existing function or assurance. Cyber defense technology and secure systems engineering must be advanced to the stage where it's a highly usable enabler for the rapid pace of new functionality instead of an impediment.

- *Metrics-Based Quantifiable Security*. Where possible, create the ability to quantitatively or even qualitatively determine the extent to which critical systems can withstand attacks based on realistic assumptions. Without such metrics, it's hard to judge progress and assess the effectiveness of proposed solutions. Metrics are fundamental.

5. Organizing for Success

5.1. The Right Organizing Models—NASA and DARPA

Strategic cyber defense to avert catastrophically damaging attacks will take careful orchestration of many government (for example, DHS, DoD, the intelligence community, Department of Treasury, Department of Energy, Department of Commerce, Department of State, and Department of Justice) and private-sector organizations (critical infrastructure providers).

A Manhattan-Project-style comprehensive program will require a high-talent, focused, agile group with presidential budgetary authority. This is the only way it will succeed. The group will have to orchestrate planning, engineering, operations, and research toward strategic national cyber defense capabilities.

To achieve the goal of the proper characteristics of a leading organization, we take inspiration from the Defense Advanced Research Project's Agency (DARPA) and the National Aeronautics and Space Administration (NASA).

DARPA is an elite team of some of the country's best researchers, leading large programs that change the world—200 program managers execute a budget of $3 billion per year, creating technology such as the Internet. DARPA has exceptional authorities such as the ability to hire top talent (using, for example, the Intergovernmental Personnel Act [IPA]) and "other transaction" authority to enter into commercial contracts for exceptional staff to do exceptional things. DARPA was created in response to Russia's launch of Sputnik, and its mission became "to avoid technological surprise." We've seen the equivalent launch of cyber space "Sputniks" from Russia and China in recent years, so the creation of a similar entity makes good sense.

NASA started out as a special projects office with presidential support. Such support became particularly strong under President Kennedy, with his goal of putting "a man on the moon by the end of the decade." Once the right people were in place, the right processes created, and the work stabilized, a bureaucracy formed to institutionalize what was working well.

5.2. The Right Start—Agile SPO

So, using these models, we need a quick-start (within 90 days of the new presidential administration's start) special projects office (SPO) of a hand-picked mix of top-talent engineers and computer scientists as well as economists, political scientists, and sociologists, to initiate the programs that will provide a national strategic cyber defense.

A SPO model is recommended because it can receive special authorities and launch quickly within an existing government infrastructure. A model that starts a new mission by creating a bureaucracy leads to a three-year-delay, as we've seen with examples such as the creation of DHS. We must start with a SPO—call it the National Cyber SPO—to work out the best ideas and processes and then institutionalize them.

5.3. The Right Place—Above the Fray

Both DARPA and NASA have strong connections to the DoD culture (although NASA's predecessor was an independent advisory committee reporting to the president), so one option would be to start the SPO in the DoD. Culturally, the DoD isn't focused on defending the homeland and it still doesn't see cyber space as a place of intrinsic primary values (money is now just bits in computers, for example); rather, it views this territory as a support element for prosecuting physical war. This culture will eventually change as more and more leaders understand the primacy of cyber space. We simply can't wait for this change to become pervasive in the DoD.

Furthermore, the execution of a major cyber initiative requires the coordination and orchestration of many agencies, none of which would want to be dominated by the other. Naming one of the agencies the lead agency wouldn't create the ideal environment.

Thus, it makes sense to create the SPO in the White House as a temporary measure of expediency.

5.4 The Right Clout—Presidential Backing

The next president should be 100% behind the program and show continued support and interest and grant it the A-1 resource priorities given to such programs as the Manhattan Project. The president should hold the initiative up as one of his crown jewels and continue to keep its vision in front of the American people and keep the staff motivated. Housing the SPO in the White House would enhance and contribute to this connection.

4.5 The Right Authority—Power of the Purse

It would be impractical and undesirable to pull out all the cyber-related organizations from myriad agencies executing a piece of the program into one institution. The existing bureaucracies are needed to effectively execute the broader national priority program. Creating a matrix management structure might be useful, but ultimately the authority to coordinate comes from the power of the purse—the authority to control budget.

So, any new money associated with executing the program should be put under the control of the SPO in that it would have approval authority over those budget items. To be clear, we aren't talking about reprogramming all cyber-related budgets from existing agencies. Existing programs required to prosecute the mission-specific portions of those agencies should continue with existing processes. New monies associated with achieving the objectives of the national initiative should be completely under the purview of the National Cyber SPO—delegated, in part to other agencies as executive agents, and retained in part and executed as programs in the National Cyber SPO.

5.6. The Right Leadership—Best of Best

Although the head of the National Cyber SPO would certainly be a political appointee who works directly for the president as a special assistant, this person would have to be the best of the best in the country—a generic manager simply won't do. This visionary must be trusted by the scientists and engineers working in the organization for his/her competence and leadership. This person must also be trusted by industry and academia alike because he/she will be an essential part of the solution. The president should pick this person with great care and sound advice and counsel from his closest advisors and from experts in the cyber realm who know the short list of people with the requisite skills and credentials.

5.7. The Right Size

The organization should be kept lean and agile—200 people are about right (using DARPA as a model). Authority for hiring directly from industry should be conferred onto the SPO immediately. If it chooses 20 of the country's best, then those people can seek out and attract the rest. Use Scientific and Engineering Technical Assistance (SETA) support liberally to amplify the capabilities of the cadre in the way that DARPA does. Give individuals complete authority over their programs and trust them to do the right thing toward the stated vision. Hire carefully—the National Cyber SPO director should be personally involved in each hire.

5.8. The Right Mission

The mission statement for this SPO should be simple and clear. As a strawman, it could be to "avoid strategic damage to the United States from Cyber space" or more positively, "make the United States the world's cyber space superpower," or less militaristically, "lead the world in cyber space safety and security for the benefit of mankind."

6. A Way Forward

Designate a government leader of a small transition team of top talent as discussed above. Give a three-month hard deadline for the team of experts to develop a "blueprint" to launch the project, including technical and program management aspects. The team would be responsible for working out the technical plan and the details of effective organization structure. An overarching architecture and roadmap is a must.

Some important "moves" that the blueprint team would need to keep in mind include the following. Using a chess analogy, we can call these "opening strategic moves" to indicate the need for immediate action and to acknowledge the broader context and dynamics of what the rest of the world is doing in cyber space. A move is a strategic action to mitigate a strategic risk in cyber defense and contributes toward a higher-level goal such as a capability or end-state. The following are quoted from the NCDI Opening Moves Workshop [4]:

- **Use Architecture Principles.** Embrace architectural principles that enable the creation and operation of secure systems. Organize networks and systems physically and logically to "operate through attacks"—so that fall-back operations and rapid recovery and repair from attacks, even of an unanticipated nature, are possible. As a policy, favor stratified/partitioned designs for critical security components. Re-organize networks that have moved away from these concepts. Separate critical data and functions of the control plane from the operational plane. Develop special-purpose security devices in critical areas to provide high-assurance protection functionality.

- **Value and Prioritize**. Design systems to satisfy critical mission requirements. Value and prioritize critical cyber infrastructure functions. As functions are automated and integrated, demand that the cost of operating without the function (e.g. its vulnerability to cyber attack) be calculated as a means of assessing its mission-criticality. Quantify recovery and rollback.

- **Validate**. Create and combine metrics-driven security analysis, simulations, and testing. Develop adequate test and analysis environments to vet theories of defense, cyber offense, new mechanisms, and operators using the best cyber strategy and tactics. Different test environments with a range of scales will be needed, and some might even need to be domain-specific. Numerous testbeds are under development, but must be significantly improved (for example, to be more usable and to provide data and tools to support experiments).

- **Create Assured Trust**. Exploit authentication and attestation mechanisms to establish trust and justify suspicion. Authentication of individuals to each other and to machines, and machines to individuals and to other machines, is required to establish trust, especially in new environments where mobility is the norm. Trustworthy identity combined with privacy-protecting mechanisms is a prerequisite for security policy enforcement and for mechanisms such as network admission control.

- **Develop human capital**. Inaugurate national competitions in secure system engineering to attract new talent and integrate academic, industry, and government efforts. Create unclassified national security research institutes with academic, private, and government players. Revamp research funding processes to encourage long-term, focused engagement in crucial areas. Increase funding in areas that will create a cyber workforce of researchers, system developers, and system administrators for commercial and government-critical systems.

- **Robust Research**. Initiate research in key technology areas. A few candidate areas include

    - practical techniques and tools for the secure composition of large-scale architectures, to support safe system design, extension, and evaluation;

    - transparent security mechanisms, to enable rather than interfere with work;

    - active automated forensics, to identify attackers and account for their actions;

    - self-healing and dynamic security, to raise the bar for attackers; and

    - system security benchmarking and assessment, to develop quantifiable metrics.

- **Preparatory Deep Analysis.** Some important strategic analysis has begun in various fora such as the NCDI workshop series and the Dark Angel analysis done by the PCD. This analysis should be validated and extended as part of the planning process to formulate a successful strategy for a strategic program. Mappings between what appear to be good initial moves and ultimate goals should be made and gaps should be identified and filled. Metrics should be established toward reaching the agreed on goals so that priorities can be established in maximum risk mitigation. Models for positively influencing markets to reengineer the cyber terrain toward the defender's benefit should be explored with economists and commercial industry leadership.

7. Conclusions

Focus on Strategic Cyber Space. There is a broad range of threats to almost every single system in cyber space, from individual home computers to business enterprise systems to the most critical systems that are the heart of civilization (power, telecommunications, and banking, to name the top three). All of these threats need mitigation to one degree or another. Talking about the full range of threats to the full range of systems is overwhelming and diffusing. By focusing on those grave threats that could cause strategic damage to the nation (and the world), we have the possibility of forming a cohesive and effective program.

Stakes Are High. But what if we don't do anything? Based on the vetted Dark Angel scenarios, we could compromise national security as we know it if we make a misstep today. Inaction isn't an option for any of us who now know these stakes and are entrusted by the people to *provide for the common defense* and protect the future.

Reengineering Cyber Space. Unlike normal physical territory, cyber space is engineered and manufactured. This is at once a serious problem and a great opportunity. Economic incentives in existing markets have driven cyber space to be highly functional, yet they're poorly assured from a national strategic threat perspective. The tragedy of the commons prevents the market, on its own, from addressing this large risk. At the same time, the government can't solve this problem on its own because the engineering and manufacture of the fabric of cyber space can't and shouldn't be taken over by government. Top talent from industry must therefore work hand-in-hand with the government to reengineer the markets themselves to significantly change the landscape of cyber space and make it inherently safer from strategic threats. The form such market reengineering takes, such as creating proper incentives, is beyond this paper's scope. The point is to argue that the citizens of cyber space must come together to address the potential tragedy of the commons we face.

Strategic Uncertainty Is Unacceptable. Some have argued that the threat stated in this paper is overblown. However, the author believes that the case made here is compelling, and that recent events in cyber space such as attacks in Estonia and Georgia and China's apparent espionage activities tend to confirm the strategic concern. Nonetheless, such arguments confuse policy makers and delay resolute action. So, if there is one thing that must be done immediately and with high priority, it is to validate the gravity and source of the strategic risk to our information infrastructure. The best of the best should be engaged in this analysis because without it, any significant program created to deal with the threat is likely to waiver and wander from its focused goal, thus endangering the program.

The nation and the world stand on the cusp of an information age whose infrastructure is a vast, new, untamed territory with much promise but whose perils must be brought under control. It's past time to begin that process.

8. References

[1] United States Constitution, Preamble.

[2] Saydjari, O.S., "Testimony of O. Sami Saydjari, President, Professionals for Cyber Defense," 25 April 2007; http://homeland.house.gov/SiteDocuments/20070425145307-82503.pdf.

[3] Danzig, R., *Catastrophic Bioterrorism—What Is to be Done*, Center for Technology and National Security Policy, Aug. 2003.

[4] NCDI, "National Cyber Defense Initiative 'Opening Moves' Workshop Report," 3-7 Dec. 2007, Monterey, California; http://ncdi.nps.edu/.